

French, R. M. & Thomas, E. (2001). The Dynamical Hypothesis in Cognitive Science: A review essay of *Mind As Motion*. *Minds and Machines*, 11, 1, 101-111.

The Dynamical Hypothesis in Cognitive Science: A review essay of *Mind As Motion*

ROBERT M. FRENCH

Department of Psychology, University of Liège, Liège, B-4000 Belgium
email: rfrench@ulg.ac.be

ELIZABETH THOMAS

Institut Léon Frédéricq, University of Liège, Liège, B-4000 Belgium
email: ethomas@ulg.ac.be

Review of: Robert F. Port and Timothy Van Gelder, *Mind As Motion*, Cambridge, MA: The MIT Press, 1995, 608 pp., 139 illus., \$80.00 (cloth), ISBN 0-262-16150-8, \$18.95, (paper), ISBN 0-262-66110-1

Introduction

Sometimes it is hard to know precisely what to think about the Dynamical Hypothesis (hereafter, DH), the new kid on the block in cognitive science and described succinctly by the slogan “cognitive agents are dynamical systems” (Van Gelder, 1998). Is the DH a radically new approach to understanding human cognition? (No.) Is it providing deep insights that traditional symbolic artificial intelligence overlooked? (Certainly.) Is it providing deep insights that recurrent connectionist models, circa 1990, had overlooked? (Probably not.) Is time, as instantiated in the DH, necessary to our understanding of cognition? (Certainly.) Is time, as instantiated in the DH, sufficient for understanding cognition? (Certainly not.)

The DH is often touted as being a revolutionary alternative to the traditional Physical Symbol System Hypothesis (PSSH, renamed in *Mind As Motion*, the Computational Hypothesis) (Newell & Simon, 1976) that was the bedrock of artificial intelligence for 25 years. We disagree.

First, as pointed out by the authors themselves, the use of dynamics as a framework to comprehend the brain and cognition is not new. In the early 1950's, for example, W. Ross Ashby wrote a wonderful little treatise called *Design for a Brain* (1952) based on the recurrent, dynamical nature of the brain. The whole field of cybernetics (Wiener, 1948), developed during the late 1940's, was also deeply concerned with feedback and stability in complex evolving systems. The explicit goal of these authors and others was to develop a general framework within which biological and cognitive systems could be studied and understood. These efforts, while visionary in scope, did not bear fruit, in part because the systems they hoped to understand — in particular, human brain function and cognitive activities — were extraordinarily complex and there were no realistic means of empirically testing hypotheses generated by this approach to cognition.

Second, the really revolutionary alternative to the PSSH — connectionism — was provided by researchers like Stephen Grossberg (1976), J. J. Hopfield (1982, 1984), Teuvo Kohonen (1982), David Ackley, Geoffrey Hinton, and Terrence Sejnowski (1985), and, above all, David Rumelhart, James McClelland and the PDP Research Group in 1986 when their book, *Parallel Distributed Processing: Explorations in the Microstruction of Cognition* (Rumelhart & McClelland, 1986), appeared. Within the context of a critique of the PSSH, the Dynamical Hypothesis represents an incremental, albeit important, extension of the work begun by the connectionists in the early- and mid-1980's. In particular, the DH's insistence on the importance of *analyzing the evolution of internal patterns of activation over time* is a major contribution.

Virtually from the outset, however, the static, stimulus-response nature of a large class of connectionist models introduced in the mid to late 1980's was recognized (see, for example, Kaplan, Weaver, & French, 1991) and led to a second-wave of *recurrent* connectionist models. For example, Jeffrey Elman's ground-breaking paper, "Finding structure in time" (Elman, 1990), in which he introduced the Simple Recurrent Network, clearly indicates the singular importance he attached to time. Once modelers began working with recurrent connectionist networks, analyzing the evolution of their internal patterns of activation over time, determining how their internal representations clustered, etc., they were fully part of the Dynamical Hypothesis paradigm.

A reasonable evaluation of the contribution to the field of *Mind as Motion* would be to say that, even though the contributors to the book were not the first to attend to the issue of time in cognitive modeling, their systematic emphasis on studying cognition as a phenomenon that evolves over time has brought the issue of time to center stage in cognitive modeling. Their application of dynamical systems tools has provided the framework with which this approach can be carried out. And therein lies the importance of this book.

The Physical Symbol System Hypothesis

The ultimate purpose of *Mind as Motion* is to persuade the reader that a dynamical systems approach to cognition provides a new and extremely useful approach to understanding cognition. And, in many respects, they argue their case admirably. This new approach is claimed to not only provide new practical tools with which to study the mind but indeed to change the very way in which we approach the question of the workings of the mind. The "dynamical" approach to mind does, indeed, lie in stark contrast to the traditional computational approach whose roots lie in formal logic and the Physical Symbol System Hypothesis of Newell and Simon (1976). The PSSH claims to provide a necessary and sufficient condition for intelligence. It is rooted in a number of explicit assumptions — namely,

- that the world can be cut up into discrete objects, each of which can be designated by a symbol;
- that each symbol refers to an object, an action, or a state of the world (e.g., an object — *dog, cat, piano, etc.* — or an action — *bark, sleep, shout, etc.*);
- that each string of symbols (i.e., an expression) has an interpretation in the world (e.g., *People hate mosquitoes.*);

- that rules and an underlying “logic of thought” govern the manipulation of the symbols and expressions in the system;
- that cognitive agents are, basically, digital computers.

In other words, according to the PSSH, if we have a rich enough repertoire of symbols and a detailed enough collection of rules for manipulating these symbols, intelligence will necessarily ensue. And not only that, *all* intelligence is necessarily a product of this type of symbol manipulation. The adherents to this point of view do not deny that symbols are engendered by lower-level processes, for that would be madness, but they do deny the importance of these lower-level processes in the study of cognition. Representations, for the adherents of the PSSH, are, for the most part, static structures made up of lists of properties for each object represented (which is one of the reasons why a list processing language like Lisp was an ideal tool to process them). One of the major problems with this approach is its disavowal of the importance of the underlying mechanisms that give rise to the symbols and rules manipulating them. It turns out that there are a great many cognitive phenomena that cannot be modeled at all when their “subcognitive” aspects are ignored.

The Dynamical Hypothesis

In contrast, in the dynamical approach to cognition, representations are no longer static structures that are unrelated to the underlying hardware of the brain. Instead, they are defined by the evolution of patterns that emerge from a substrate of underlying hardware and that are detected in the geometry of time.

Compare how a traditional AI system and a dynamical system might solve the problem of binding. Consider what happens when we recognize a red car. We know that color and shape are processed in different areas of the cortex. Consequently, the brain must somehow conjoin the concept “red” and the concept “car” in order for us to perceive a “red car.” We say that the concept “red” is *bound* to the concept “car.” Now a traditional symbol system simply makes a copy of both concepts and joins the two concepts in Working Memory. If the system needs to make more copies of either concept, it does so. There are no theoretical constraints on this copying processing that transfers concepts from long-term memory to working memory.

Now consider a dynamical-system way of doing binding that uses the brain’s 40 Hz synchronous oscillations. A red car is not represented by a neuron (or even a cluster of neurons) that responds to both “car” and “red” or by a direct anatomical connection between the a neuron for “red” and another for “car.” Rather, the binding of these disparate concepts instead takes place *in time*, by the synchronous firing of these separate clusters of neurons, one corresponding to the concept “red” and another corresponding to the concept “ball” (Shastri & Ajjanagadde, 1993; Sougné & French, 1997; Sougné, 2000).

According to the Dynamical Hypothesis, it is crucial to monitor activity as it unfolds in time (time series). Sometimes the time series creates a pattern in space called an attractor. All of cognition, according to the Dynamical Hypothesis, takes place via transitions between these attractors. Unlike the PSSH, where the transitions between fixed entities called symbols are discrete and generally instantaneous, transitions between

attractors in the dynamical hypothesis paradigm take place over time and therefore intrinsically incorporate the various transitional states that may occur before settling to a final state.

Introducing the Dynamical Hypothesis

The book begins with an excellent introduction by the editors, Timothy van Gelder and Robert Port, appropriately entitled “It’s About Time: An Overview of the Dynamical Approach to Cognition,” in which the overarching philosophy of the dynamical approach to cognition is presented. In terms of its explanatory value, this chapter is the high-water mark of the book, at least for non-specialists. This first chapter includes a detailed description of the inadequacies of the PSSH. Most of these criticisms can be found, in one form or another, in the early connectionist attacks on traditional artificial intelligence of the mid-1980’s. The contribution of the Dynamical Hypothesis is to shift the focus of the connectionist criticisms of the PSSH. The connectionists focused on the advantages of distributed systems with massive parallel processing of information. The dynamical hypothesis focuses, in addition, on the evolution of the state of the system (or parts of the system) over time.

When dynamicists talk about “time,” they are really talking about sequences of events in time, whether these events are at the level of neural firings, arm movements of babies, or high-level actions. And in this sense time clearly was an issue in many symbolic systems — planning, for example, was based on performing a *sequence of actions* in order to achieve a goal, which was a time-dependent process. Thus, there are numerous examples of researchers in the PSSH tradition studying sequences of high-level actions and their evolution over time. For dynamical hypothesists, actions that occur over time, whether they are at the neural, cognitive, or even social level, are the central focus of their work. In addition, one of their main goals is to understand how higher level cognitive behaviors might arise from sequences of lower level actions. So, for example, for a dynamicist whose focus is on sequences of neural firings, the goal would be to understand how low-level neural activity might give rise to our perception of high-level concepts. In contrast, the PSSH is wholly unconcerned with this lower level. This undoubtedly marks one of the most fundamental differences between the two approaches to modeling cognition. Researchers in the PSSH tradition are explicitly uninterested in linking the neural level to the higher symbolic level. Dynamicists, on other hand, would argue that understanding cognition requires an understanding of how temporal sequences of low-level events produce the phenomena (e.g., the perception of categories, relations between categories, etc.) that are finally perceived by us.

Unfortunately, after the first introductory chapter, *Mind As Motion*, reads more like a thinly veiled proceedings of a technical conference on dynamics than “the first comprehensive presentation of the dynamical approach to cognition” that it is touted to be. Despite considerable efforts on the part of the editors, they nonetheless failed to ensure that the individual chapters were written in a manner that would make them accessible to cognitive scientists without considerable mathematical training. The technical difficulty of the content has undoubtedly meant that the book has not had its hoped-for impact. There is a clear conflict between the desires of the editors (“presenting

real research papers”) and the needs of the cognitive science research community interested in an general understanding of the dynamical-systems approach to cognition.

This is the first significant difficulty with *Mind As Motion*. The second is that the authors, in seeking to apply the tools of dynamical systems, often analyze physical quantities that one has difficulty identifying with cognition.

The first difficulty is exemplified by the second chapter by Alec Norton (“Dynamics: An Introduction”), which is supposed to be a short introduction to dynamics for people not familiar with the mathematical underpinnings of the subject. After a few elementary definitions, the author moves rapidly into the technical language of differential equations and dynamical systems theory and the chapter ceases to be useful to the uninitiated. And the five-page glossary added at the end of the book does little to improve the situation. For example, here is how the glossary “explains” the concepts of *diffeomorphism*: “This is a differentiable mapping of a space to itself with a differentiable inverse. For example, the time-one map of a flow is a diffeomorphism.” and *manifold*: “A geometric object that has fewer dimensions locally than the parametric space in which it is embedded....Generally, a manifold is a topological space that is locally homeomorphic to \mathbb{R}^n , and is the most common setting for much of geometric dynamical systems theory.” And the explanation of key concepts like *flow*, *gradient system*, *orbit*, *Poincaré map*, *trajectory*, and *vector field* are equally technical for non mathematicians. The point is that either one has a relatively strong background in mathematics, in general, and in differential equations, in particular, and doesn’t need the chapter or, one doesn’t have such a background and the chapter is of little use.

Tools for the study of dynamical systems

As we have already said, one of the major contributions of *Mind as Motion* is advocating studying cognitive processes with a new arsenal of tools, largely developed in the domain of the physics of complex, dynamical systems.

One type of dynamical systems approach to understanding a particular phenomenon is to attempt to develop a set of equations that describe the evolution over time of the entire system. Marco Guinti calls these *Galilean dynamical models* (chapter 18, “Dynamical Models of Cognition”). This technique, one that has long been used fruitfully in physics and engineering, is essentially as follows. A variable of the system which is deemed to be important (a “cognitive magnitude”) is identified and an equation is formulated to describe the evolution of this magnitude over time. It is also common to use a set of dynamical equations to characterize various *components* of a system rather than the system as a whole. Examples of this approach are found in the various connectionist models described in the book. Equations are used to describe various nodes in the network without an attempt to introduce equations that describe the overall behavior of the network.

The ultimate goal of dynamicists in cognitive science is, presumably, to develop precise sets of equations to describe the behavior of real cognitive systems. This, however, may never be possible due to the difficulties inherent in actually developing equations that precisely describe the dynamics of such a complex system. A more realistic approach to understanding the dynamical behavior of cognitive systems is to make use of the graphical techniques — in particular, *phase portraits* — currently available

Losing sight of cognition

We now turn to some of the more philosophical aspects of the dynamical hypothesis for cognitive systems. First, we will consider the question of whether the dynamical systems framework replaces the PSSH. As we pointed out above, one of the major problems with the PSSH approach is that it places no importance on the underlying mechanisms that give rise to symbols and the relationships between them. But many cognitive phenomena, such as priming, color-contour binding, etc., are virtually impossible to model if “subcognitive” (and neural) aspects are ignored altogether. Donald Hebb (1949) was one of the first to explicitly address the issue of how low-level modifications of neurons gave rise to high-level cognitive phenomena, in particular, learning. The rise of connectionism in the 1980’s was an explicit response to the subcognitive shortcomings of traditional artificial intelligence, just as the current interest in the Dynamical Hypothesis is an extension of the overall concerns of connectionism. Many of these issues have been discussed at length in review articles on the dynamical hypothesis by van Gelder (1998) and Randall Beer (2000).

The dynamical hypothesis can be considered to be an extension of an information processing approach to cognitive science, but with a significantly different focus than the traditional artificial intelligence information processing paradigm. One of the dangers, however, of the headlong rush to apply dynamics to all levels of cognitive processing is losing sight of cognition. Understanding the deeper meaning of period doubling in a dynamical system, for example, requires maintaining a clear focus on the high-level task to be accomplished and the means by which the system can accomplish it. While it may seem unnecessarily alarmist to speak out against the proliferation of such research, one need not look very far to see myriad studies of dynamics in which the higher level products of the dynamics are barely mentioned. Neuroscience journals are full of complex descriptions of changes in the dynamics of neurons and neural systems that never touch base with the purpose of these changes. They could just as easily be studying the dynamics of digestion.

Consider, for example, Esther Thelen’s chapter on the development of embodied cognition in infants (chapter 3, “Time-Scale Dynamics and the Development of an Embodied Cognition”). The author goes to considerable length to convince us of the importance of the dynamics of a child’s limb movement in the emergence of a higher level awareness of its surroundings. Thelen explicitly links the dynamics of reaching to emergent intentionality. While we are convinced of the ultimate importance of embodiment (i.e., the necessary interaction with the environment of the body in the development of cognition), her explicit link between reaching and intentionality is far from obvious. One could do a thorough dynamical systems analysis of an amoeba’s pseudopod extension towards a food particle, but this would certainly have nothing to do with emerging intentionality. So one key question, which goes largely unanswered in this book, is what types of dynamical systems are related to cognition, what types are not, and why. (Dennett, 1998, makes a related point.)

The importance of embodiment

Let us focus for a moment on the notion of embodiment in cognition. This idea, of considerable importance in the dynamical hypothesis framework, is that it is impossible to understand cognition without reference to the bodies in which the cognitive processes take place (Lakoff, 1986; Varela, Thompson, Rosch, 1993; etc.). In other words, the disembodied, brain-in-a-vat cognition that has long been the plaything of philosophers of mind is, in this view, nonsensical. Minds are products of bodies, of their structure, their capacities and their limitations. And, as a result, it makes sense to study the dynamics, not just of neurons, but of bodily movements in an attempt to study cognition. This is what dynamicists mean when they talk about closing the dualistic gap between the mind and the body, even though at the present stage of our understanding, we believe that it is more reasonable to talk about narrowing, rather than closing, this gap.

In the traditional PSSH framework, cognition is the manipulation of various symbols or representations in the head. But the low-level mechanisms that actually produce transitions from one representation to another are generally absent from these models. Thus, another advantage of a dynamical systems approach is that it provides a parsimonious explanation for these transitions between these various representations. From a dynamical systems perspective, the states of the system emerge from the dynamics of the system itself. By varying certain critical parameters, the system makes a transition (bifurcation) to a new state. Jordan Pollack (chapter 10, "The Induction of Dynamic Recognizers"), for example, shows how his network is able to perform a serial parity task only after a small change in the weights of the network produces a radical change in the behavior of the network. The network's limit behavior changes from being a limit point to a limit cycle. Similarly, Jean Petitot's article (chapter 9, "Morphodynamics and Attractor Syntax: Constituency in Visual Perception and Cognitive Grammar") on language processing models the individual terms of a sentence by attractors and the bifurcations between these attractors as the syntax relating the attractors to one another.

This type of system is said to be 'self-organizing' and eliminates the need for the "central executive" that is so crucial to traditional AI systems. The greater parsimony entailed by the elimination of the need for a global executive to orchestrate the appropriate state transitions of the system is one of the things that makes dynamical systems approaches to cognitive science so appealing.

Conclusion

This is a book about the importance of the evolution of time in modeling psychological processes. It presents various ways in which studies of cognition should incorporate time, among them the use of equations to describe changes of cognitive magnitudes over time, the use of recurrent connectionist models, or the use of graphical methods to uncover structure in the temporal evolution of a process. It calls for the incorporation of the framework of dynamics, which deals specifically with the temporal evolution of a system, into the study of cognition. While this is certainly a promising approach, many problems remain. These problems are too infrequently discussed and, as a result, the uninformed reader is left with an inaccurate impression of the vast difficulties that lay ahead. For example, much of current dynamical theory has been developed for physical systems which do not show the problems of noise and lack of stationarity displayed by

biological and other complex systems (Crutchfield, 1998; French & Thomas, 1998; Rapp 1993). Further, the characterization of attractors requires huge amounts of data, implying correspondingly huge amounts of calculation. Since in real cognition, the transition from one state to another is extremely rapid, one must wonder how the types of highly computationally intensive attractors described in this book could be processed by a real brain. While certain authors (e.g., Thelen and Reidbord) do mention their inability to characterize an attractor due to insufficient data, this crucial point is almost overlooked entirely in this book. There should at least have been references to the methods in nonlinear dynamics that have been specially developed to deal with short time series.

While the ideas presented in this book are important, it does not live up to its dust-jacket description as the first comprehensive presentation of the dynamical approach to cognition. Unfortunately, *Mind As Motion* comes up short in the difficult task of conveying the ideas of dynamics as they apply to cognition to those readers not already familiar with them.

Acknowledgments

The work was supported in part by research grants ARC 99/044-241 and IUPA 4/19 from the Belgian government and HPRN-CT-1999-00065 from the European Commission.

References

- Ackley, D., Hinton, G., Sejnowski, T. (1985). A learning algorithm for Boltzmann machines. *Cognitive Science*, 9: 147-169.
- Ashby, W. R. (1952). *Design for a Brain*. London: Chapman and Hall.
- Beer, R. (2000). Dynamical approaches to cognitive science. *Trends in Cognitive Sciences*, 4(3). 91-99.
- Crutchfield, J. (1998). Dynamical embodiments of computation in cognitive processes. *The Behavior and Brain Sciences*, 21, 635.
- Dennett, D. (1998). Revolution, no! Reform, si! *Behavior and Brain Sciences*, 21, 636-637.
- Elman, J. (1990). Finding structure in time. *Cognitive Science*, 14, 179-211.
- French, R. & Thomas, E. (1998). The Dynamical Hypothesis: One Battle Behind. In *The Behavior and Brain Sciences*, 21:5, 640-641.
- Grossberg, S. (1976). Adaptive pattern classification and universal recoding: I. Parallel development and coding of neural feature detectors. *Biological Cybernetics*, 23: 121-134.
- Hebb, D.O. (1949). *The Organization of Behavior*. NY, NY: Wiley.
- Hopfield, J. (1982). Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences*, 79: 2554-2558.
- Hopfield, J. (1984). Neurons with graded response have collective computational properties like those of two-state neurons. *Proceedings of the National Academy of Sciences*, 81: 3088-3092.

- Kaplan, S., Weaver, M., and French, R. (1991). Active Symbols and Internal Models: Towards a Cognitive Connectionism. *AI and Society*, 4(1), 51-71. Reprinted in A. Clark & R. Lutz (eds.), *Connectionism in Context*. Springer-Verlag, 1992, pp. 91-110.
- Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological Cybernetics*, 43: 59-69.
- Lakoff, G., (1986). *Women, Fire and Dangerous Things*. Cambridge, MA: MIT Press
- Minsky, M. & Papert, S. (1969). *Perceptrons*. Cambridge, MA: MIT Press.
- Newell, A. & Simon, H. (1976). Computer science as empirical inquiry: Symbols and search. *Communications of the ACM*, 19, 113-126.
- Rapp P. (1993) Chaos in the neurosciences: cautionary tales from the frontier. *Biologist*, 40, 89-94.
- Rosenblatt, F. (1962). *Principles of Neurodynamics: Perceptron and the Theory of Brain Mechanisms*. Washington, D.C.: Spartan Books.
- Rumelhart, D. & McClelland, J. (1986). *Parallel Distributed Processing*. Cambridge, MA: MIT Press.
- Shastri, L. and Ajjanagadde, V. (1993). From simple associations to systematic reasoning: A connectionist representation of rules, variables, and dynamic bindings using temporal synchrony. *Behavioral and Brain Sciences*, 16, 417-494.
- Sougné, J. & French, R. (1997). A Neurobiologically Inspired Model of Working Memory Based on Neuronal Synchrony and Rythmicity. In J. Bullinaria, G. Houghton, D. Glasspool (eds.). *Connectionist Representations: Proceedings of the Fourth Neural Computation and Psychology Workshop*. Springer-Verlag. 155-167.
- Sougné, J. (2000). Binding and Multiple Instantiation in a Distributed Network of Spiking Nodes. *Connection Science* (in press).
- Takens F.(1981) Detecting strange attractors in turbulence. In *Dynamical Systems and Turbulence*, D.A. Rand and L.-S. Young, eds. Springer-Verlag, New York, pp. 366-381.
- Van Gelder, T. (1998). The Dynamical Hypothesis in Cognitive Science. *The Behavior and Brain Sciences*, 21, 615-665.
- Varela, F., Thompson, E., & Rosch, E., (1993), *The Embodied Mind*. Cambridge, MA: MIT Press.
- Wiener, N. (1948). *Cybernetics*. New York: Wiley.